

Farm Remote Copy Utility (FCP)

What Is It for ?

How It Works ?

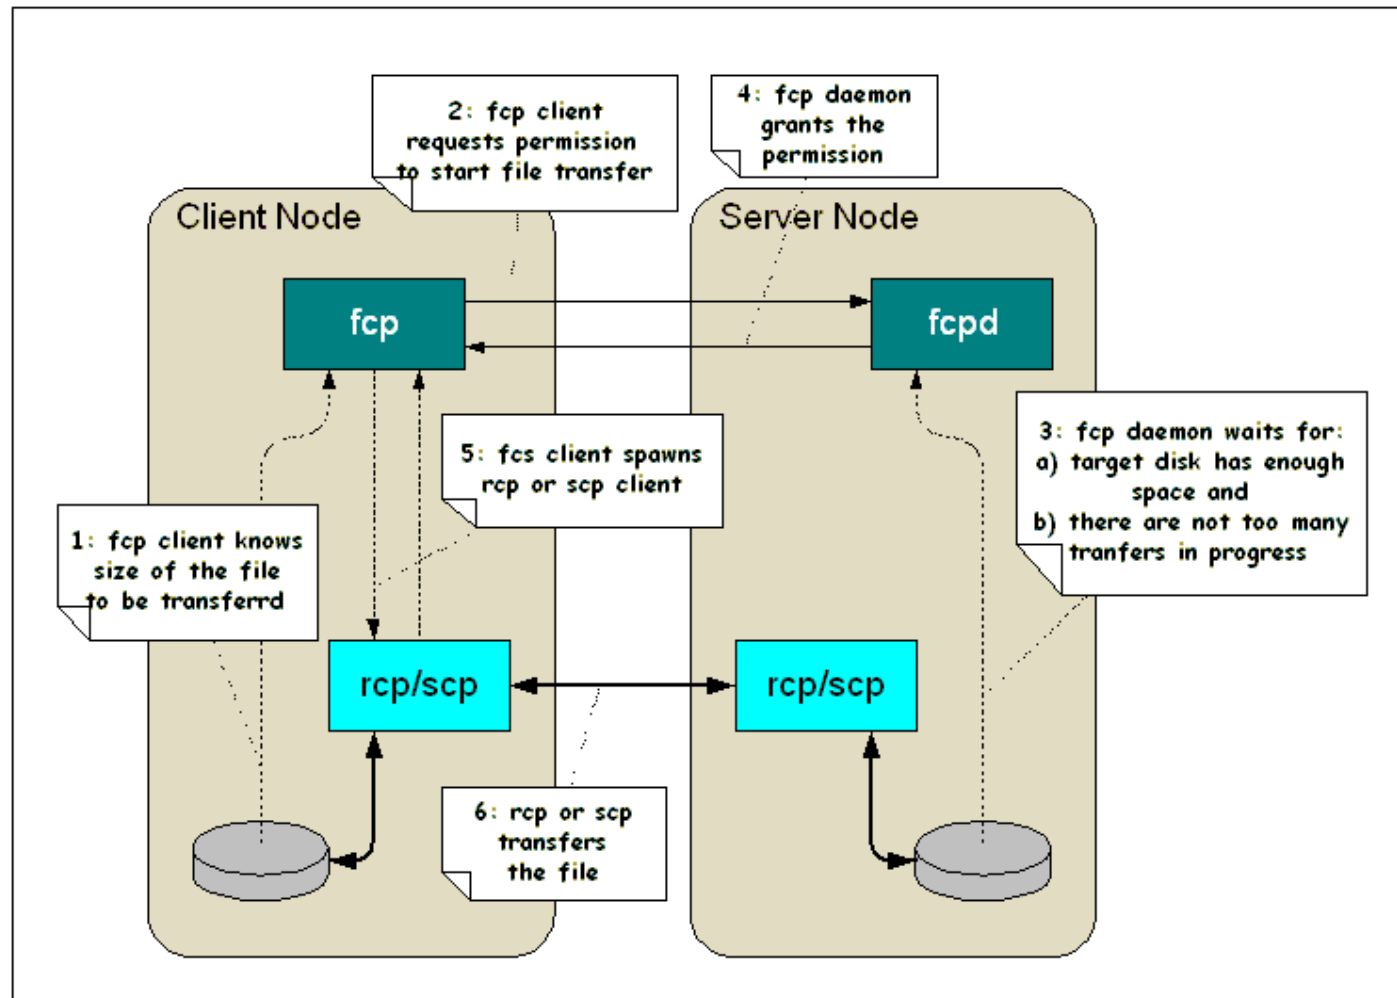
User Interface

Current Status and Future

What Is It for ?

- ◆ To help batch processes transfer data to and from permanent disk storage (I/O node)
- ◆ Potential problems to solve when transferring data between the I/O and a worker node:
 - Limited network bandwidth
 - Limited through-put of the I/O node
 - Limited size of the upload target disk
- ◆ What FCP does:
 - Limits number of concurrent data transfers
 - Does not allow to start an upload until there *is* and *will be* enough disk space to finish it

How It Works



User Interface

◆ Basic functions:

■ UPS setup:

```
$ setup fcp
```

■ Copying a file using rcp:

```
$ fcp /path/data.file server:/path/file
```

```
$ fcp server:/path/file /path/data.file
```

■ Specifying rcp options:

```
$ fcp -c rcp -p server:/path/file /path/data.file
```

■ Using scp:

```
$ fcp -c scp -c none data.file server:/path/file
```

■ Wildcards are not supported (yet):

```
$ fcp /tmp/* server:/tmp      # will not work !
```

■ Verbose mode:

```
$ fcp -v ~/.cshrc fnsfo:/tmp/.cshrc
```

```
Request upload...
```

```
Request granted.
```

```
Spawning <rcp /home/farms/.cshrc fnsfo:/tmp/.cshrc> ...
```

```
Transfer status = 0
```

User Interface

◆ Get server status

```
$ fcps fnsfo
```

RID	user.pid@host	TS	Free/Avbl (MB)	Size	File
600	farms.53429604@fnsfo	U*	3287/2981	101	/tmp/file.3
601	farms.57052704@fnsfo	U*	3287/2981	100	/tmp/file.4
602	farms.55467993@fnsfo	U*	3287/2981	105	/tmp/file.5
603	farms.55397610@fnsfo	D*		99	/tmp/file.1
604	farms.36191@fnpc230	D		93	/tmp/file.1
605	farms.25911@fnpc221	D		85	/tmp/file.1

◆ Aborting a request

- If you know the transfer has finished:

```
$ fcp_abort 601
```

- Transfer will be automatically aborted after configurable time-out (say, 30 minutes)

Current Status and Possible Future

- ◆ Installed on fnsfo farm, beta-tested and evaluated by E781
- ◆ Basis for possible future development
 - *Reservations*: I want to copy 10 files 500 MB in total, reserve 500 MB for me on the target disk
 - *Multi-disk placement*: I have 5 disks, choose whichever is available, put my file there and let me know where it is stored
 - *Improved scheduling*:
 - Currently, simple queue discipline
 - Small transfers deserve higher priority
 - Big transfers should not wait indefinitely
 - Prioritization on per-user, -group, -project basis
 - *Improved reliability*: Do not use rcp/scp, use own transportation means – know exactly when transfer is over
 - Note: security issues will have to be addressed